



Docket No.: 500.43371X00

IFW

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re the Application of:

Shougo MIKAMI

Serial No. 10/743,727

Filed: December 24, 2003

For: METHOD, APPARATUS AND COMPUTER READABLE MEDIUM
FOR MANAGING REPLICATION OF BACK-UP OBJECT

SUPPLEMENTAL PETITION TO MAKE SPECIAL
UNDER 37 CFR §1.102(MPEP §708.02)

June 13, 2005

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Supplemental to the Petition to Make Special filed on May 16, 2005,

Applicants submit the following additional remarks.

It is submitted that the cited references, whether considered alone or in combination, fail to disclose or suggest the invention as claimed. In particular, the cited references, at a minimum, fail to disclose or suggest in combination with the other limitations recited in the claims:

a first feature of the present invention as recited in independent claim 1 including holding a first table that, if files as backup objects are classified into plural groups, stores each group name, file names, and each policy in a manner of being made to correspond to each other, said file names being included in said each group, said each policy being a method for managing replication data, and

when transmitting said pair information to said storage apparatus in response to a receipt of a specified second group name, said second group including said first file, said pair information including said first original-volume name and a second duplicate-volume name, said first original-volume storing therein the data of said first file, extracting said first duplicate-volume name out of said second table, said first duplicate-volume storing therein replication data of said first file; and

a second feature of the present invention as recited in independent claim 2 wherein said computer holding a first table and a second table if files as backup objects are classified into plural groups, said first table being designed for storing file names in a manner of being made to correspond to each group name, said file names being included in said each group, said second table being designed for storing said file names and said group names in a manner of being made to correspond to said duplicate-volume names, and said replication-data management method comprising a step of: when transmitting said pair information to said storage apparatus in response to the receipt of a specified first group name, said pair information including a first original-volume name and a first duplicate-volume name, said first original-volume storing therein a first file, storing said first file name and said specified first group name into said second table in a manner of being made to correspond to said first duplicate-volume name. Independent claim 8 recites a memory unit for holding a first table and a second table when files becoming backup targets are classified into plural groups, said first table being designed for storing file names in a manner of being

made to correspond to each group name, said file names being included in said each group, said second table being designed for storing said file names and said group names in a manner of being made to correspond to said duplicate-volume names, and a volume setting unit for executing a step of: when transmitting said pair information to said storage apparatus in response to a receipt of a specified first group name, said pair information including a first original-volume name and a first duplicate-volume name, said first original-volume storing therein a first file, storing said first file name and said first group name whose specification has been received into said second table in a manner of being made to correspond to said first duplicate-volume name.

To the extent applicable to the present Petition, Applicants submit that although the distinguishing feature(s) may represent a substantial portion of the claimed invention, the claimed invention including said feature(s) and their inter-operation provides a novel storage system and system and method related to or implemented in or by said storage system not taught or suggested by any of the references of record.

The references considered most closely related to the claimed invention are briefly discussed below:

U.S. Patent No. 5,845,318 (Rose et al.) discloses a method and system for improving performance of a DASD in a data processing system by implementing a cache replacement policy which replaces cache data determined to be cheapest to retrieve from the DASD given the location of the DASD's R/W heads and updates data on a DASD from the cache when it is inexpensive in

terms of data access time or costs. The method includes utilizing a replacement policy to overwrite a selected buffer with requested data. (See, e.g., Abstract and column 2, lines 30-60, and column 5, lines 25-30.) However, unlike the present invention, Rose et al. does not disclose, at a minimum, sequential procedures of: holding a first table that, if files as backup objects are classified into plural groups, stores each group name file names, and each policy in a manner of being made to correspond to each other, said file names being included in said each group, said each policy being a method for managing replication data. Furthermore, Rose et al. do not disclose when transmitting pair information to a storage apparatus in response to a receipt of a specified second group name, a second group including a first file, a pair information including a first original-volume name and a second duplicate-volume name, a first original-volume storing therein the data of said first file, extracting a first duplicate-volume name out of a second table, said first duplicate-volume storing therein replication data of said first file. More particularly, Rose et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, and the above described second feature of the present invention as recited in independent claim 2, in combination with the other limitations recited in each of the independent claims.

U.S. Patent No. 6,131,148 (West et al.) discloses a system and method for snapshot copying of a volume of a data-storage system which employs remote dual copying techniques between a primary storage subsystem and a remote secondary storage subsystem and, more specifically, to the performance

of an equivalent of a snapshot copy of a volume in a primary storage system at a remote secondary data-storage system. The primary subsystem relays a snapshot request to a secondary subsystem, and performs a snapshot copy operation on the secondary volume. The snapshot copy of the secondary volume is then stored in a tertiary volume locally or at a remote location, in virtual memory to be later archived or to any other storage location other than the secondary volume. The current snapshot copy may overwrite a previous snapshot copy to minimize the required storage space for snapshot copies. (See, e.g., Abstract and column 8, lines 8 - 17.) However, unlike the present invention, West et al. does not disclose, at a minimum, sequential procedures of: holding a first table that, if files as backup objects are classified into plural groups, stores each group name file names, and each policy in a manner of being made to correspond to each other, said file names being included in said each group, said each policy being a method for managing replication data. Furthermore, West et al. does not disclose when transmitting pair information to a storage apparatus in response to a receipt of a specified second group name, a second group including a first file, a pair information including a first original-volume name and a second duplicate-volume name, a first original-volume storing therein the data of said first file, extracting a first duplicate-volume name out of a second table, said first duplicate-volume storing therein replication data of said first file. More particularly, West et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, and the above described second feature of the present invention as

recited in independent claim 2, in combination with the other limitations recited in each of the independent claims.

U.S. Patent No. 6,173,377 (Yanai et al.) discloses a system and method for automatically providing and maintaining a copy or mirror of data stored at a location geographically remote from the main or primary data storage device. Two data storage systems are interconnected by a data link for remote mirroring of data. Each volume of data is configured as local, primary in a remotely mirrored volume pair, or secondary in a remotely mirrored volume pair. Normally, a host computer directly accesses a primary volume, and data written to the primary volume is automatically sent over the link to a corresponding secondary volume. Each remotely mirrored volume pair can operate in a selected synchronization mode including synchronous, semi-synchronous, adaptive copy-remote write pending, and adaptive copy-disk. Direct write access to a secondary volume is denied if a "sync required" attribute is set for the volume and the volume is not synchronized. If a "volume domino" mode is enabled for a remotely mirrored volume pair, access to a volume of the pair is denied when the other volume is inaccessible. In a "links domino" mode, access to all remotely mirrored volumes is denied when remote mirroring is disrupted by an all-links failure. The domino modes can be used to initiate application-based recovery, for example, recovering a secondary data file using a secondary log file. In an active migration mode, host processing of a primary volume is concurrent with migration to a secondary volume. In an overwrite cache mode, remote write-pending data in cache can be overwritten. Write data for an entire

host channel command word chain is bundled in one link transmission. (See, e.g., Abstract and column 2, line 32, through column 3, line 21.) However, unlike the present invention, Yanai et al. does not disclose, at a minimum, sequential procedures of: holding a first table that, if files as backup objects are classified into plural groups, stores each group name file names, and each policy in a manner of being made to correspond to each other, said file names being included in said each group, said each policy being a method for managing replication data. Furthermore, Yanai et al. does not disclose when transmitting pair information to a storage apparatus in response to a receipt of a specified second group name, a second group including a first file, a pair information including a first original-volume name and a second duplicate-volume name, a first original-volume storing therein the data of said first file, extracting a first duplicate-volume name out of a second table, said first duplicate-volume storing therein replication data of said first file. More particularly, Yanai et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, and the above described second feature of the present invention as recited in independent claim 2, in combination with the other limitations recited in each of the independent claims.

U.S. Patent No. 6,304,948 (Motoyama et al.) discloses a method and system for storing and maintaining data that ensures file integrity and security in a storage system. An approach for storing and maintaining data involves determining whether data has been stored on a first non-volatile storage for a specified period of time. If so, then the data on the first non-volatile storage is

deleted by overwriting the data on the first non-volatile storage with a predetermined value so that the data cannot be recovered. The first non-volatile storage is registered with a registration authority to provide authentication of the data maintained on the first non-volatile storage. A duplicate copy of the data is stored on a second non-volatile storage and if a determination is made that the data has been stored on a first non-volatile storage for a specified period of time, then the duplicate copy of the data is deleted by overwriting the duplicate copy of the data on the second non-volatile storage with the predetermined value. (See, e.g., Abstract and column 2, lines 2-23.) However, unlike the present invention, Motoyama et al. does not disclose, at a minimum, sequential procedures of: holding a first table that, if files as backup objects are classified into plural groups, stores each group name file names, and each policy in a manner of being made to correspond to each other, said file names being included in said each group, said each policy being a method for managing replication data. Furthermore, Motoyama et al. does not disclose when transmitting pair information to a storage apparatus in response to a receipt of a specified second group name, a second group including a first file, a pair information including a first original-volume name and a second duplicate-volume name, a first original-volume storing therein the data of said first file, extracting a first duplicate-volume name out of a second table, said first duplicate-volume storing therein replication data of said first file. More particularly, Motoyama et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, and the above described second feature of the present

invention as recited in independent claim 2, in combination with the other limitations recited in each of the independent claims.

U.S. Patent No. 6,834,324 (Wood) discloses a system that emulates a tape cartridge mounted in a tape drive. The emulated tape drive communicates with a client using an interface protocol defining an addressable range. The emulated tape cartridge has a virtual addressable range that matches or exceeds the addressable range defined in the interface protocol. One or more storage appliances provide the emulation of the tape drive and a plurality of virtual volumes. An interface manager disposed between the client and the storage appliances mounts the virtual volumes as necessary to assemble the emulated tape cartridge. The interface manager has an overwrite protection feature that protects existing data from being overwritten with new data having the same block address. (See, e.g., Abstract and column 7, lines 11-34.) However, unlike present invention, Wood does not disclose, at a minimum, sequential procedures of: holding a first table that, if files as backup objects are classified into plural groups, stores each group name file names, and each policy in a manner of being made to correspond to each other, said file names being included in said each group, said each policy being a method for managing replication data. Furthermore, Wood does not disclose when transmitting pair information to a storage apparatus in response to a receipt of a specified second group name, a second group including a first file, a pair information including a first original-volume name and a second duplicate-volume name, a first original-volume storing therein the data of said first file, extracting a first duplicate-volume name

out of a second table, said first duplicate-volume storing therein replication data of said first file. More particularly, Wood does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, and the above described second feature of the present invention as recited in independent claim 2, in combination with the other limitations recited in each of the independent claims.

U.S. Patent No. 6,845,434 (Lin) discloses a method and system for updating parametric data for use in data management system. A data management system of a client device includes a data storage device for storing various parameters for configuring the data management system. The data storage device includes a backup data storage device for storing provisioning data, a random access memory device for use in dynamical provisioning data modification, and a transaction data storage device for recording variations in parameters. When the requirements of the client device are changed, the related parameters stored in the RAM device will be modified. When the modification is done in the RAM device, the data management system writes the changes to the transaction data storage device. When the data management system restarts, the system writes the data stored in the backup data storage device to the RAM device, and then updates the related parameters stored in the RAM device according to the contents of the transaction data storage device. In this way, the contents of the RAM device can correspond to the most recent requirements of the user. When the amount of content stored in the transaction data storage device has increased to a certain level, the backup data storage device will be

updated according to the contents of the transaction data storage device. Once the backup data storage device is updated with the contents of the transaction data storage device, the contents of the transaction data storage device will be erased and the transaction data storage device can store subsequent variations in parameters. The method also includes the step of being able to select a backup data storage device from two backup data storage devices according to their respective flags, and then overwriting the contents of the selected backup data storage device. (See, e.g., Abstract, column 3, line 19, through column 4, line 18, and column 9, line 50, through column 10, line 10.) However, unlike the present invention, Lin does not disclose, at a minimum, sequential procedures of: holding a first table that, if files as backup objects are classified into plural groups, stores each group name file names, and each policy in a manner of being made to correspond to each other, said file names being included in said each group, said each policy being a method for managing replication data. Furthermore, Lin does not disclose when transmitting pair information to a storage apparatus in response to a receipt of a specified second group name, a second group including a first file, a pair information including a first original-volume name and a second duplicate-volume name, a first original-volume storing therein the data of said first file, extracting a first duplicate-volume name out of a second table, said first duplicate-volume storing therein replication data of said first file. More particularly, Lin does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, and the above described second feature of the present invention as recited in

independent claim 2, in combination with the other limitations recited in each of the independent claims.

U.S. Patent No. 6,851,015 (Akahane et al.) discloses a system and method for overwriting data in nonvolatile memory and a control apparatus for using the method. Specifically, the invention provides a method of overwriting pre-overwrite data, which is the data to be overwritten, with post-overwrite data, which is the new data that overwrites the pre-overwrite data. In a nonvolatile memory such as flash memory where data is stored sector by sector, a method of overwriting a data sector is provided. The old data to be overwritten in a data sector along with its error detection code are initially saved to a backup region which is also in a non-volatile area. The old data in the data sector is then erased and new data along with its own error detection code are written into the same data sector where the old data was previously stored. (See, e.g., Abstract and column 2, lines 8 - 67.) However, unlike the present invention, Akahane et al. does not disclose, at a minimum, sequential procedures of: holding a first table that, if files as backup objects are classified into plural groups, stores each group name file names, and each policy in a manner of being made to correspond to each other, said file names being included in said each group, said each policy being a method for managing replication data. Furthermore, Akahane et al. does not disclose when transmitting pair information to a storage apparatus in response to a receipt of a specified second group name, a second group including a first file, a pair information including a first original-volume name and a second duplicate-volume name, a first original-volume storing

therein the data of said first file, extracting a first duplicate-volume name out of a second table, said first duplicate-volume storing therein replication data of said first file. More particularly, Akahane et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, and the above described second feature of the present invention as recited in independent claim 2, in combination with the other limitations recited in each of the independent claims.

U.S. Patent Publication No. 2002/0103980 (Crockett et al.) discloses a system and method for releasing storage space in a storage system where updates to a primary storage device are shadowed in a secondary storage device, in order to release storage space in the first and second storage devices. Updates to the first storage device are copied to the second storage device to provide secondary storage for the updates. A first and second tables map data sets to addresses in the first and second storage devices, respectively. A first command is detected to invalidate data sets in the first table. The addresses in the first table comprise virtual addresses, and a third table provides a mapping of the virtual addresses to physical storage locations in the first storage device. A second command is generated to update the second table to invalidate the data sets in the second storage device invalidated in the first table by the first command. A third command is detected to invalidate the virtual addresses in the third table used by the data sets invalidated in the first table to free the physical storage locations in the first storage device pointed to by the invalidated virtual addresses. A fourth command is generated that is directed to the physical

storage locations in the second storage device used by the invalidated data sets. The method also includes being able to set a flag which indicates whether a secure erase function is enabled, to allow the overwriting of data in a storage device. (See, e.g., Abstract and paragraphs 13-16.) However, unlike the present invention, Crockett et al. does not disclose, at a minimum, when transmitting pair information to a storage apparatus in response to a receipt of a specified first group name, said pair information including a first original-volume name and a first duplicate-volume name, said first original-volume storing therein a first file, storing said first file name and said specified first group name into a second table in a manner of being made to correspond to said first duplicate-volume; when transmitting said pair information to said storage apparatus in response to a receipt of a specified second group name, said second group including said first file, said pair information including said first original-volume name and a second duplicate-volume name, said first original-volume storing therein the data of said first file, extracting said first duplicate-volume name out of said second table, said first duplicate-volume storing therein replication data of said first file. More particularly, Crockett et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, and the above described second feature of the present invention as recited in independent claim 2, in combination with the other limitations recited in each of the independent claims.

U.S. Patent Publication No. 2003/0196036 (Gibble et al.) discloses a method and system for guaranteeing overwrite of expired data in a media device,

the data being organized according to logical units provided on media devices organized as physical volumes. The method comprises the steps of: maintaining information associated with a status of a physical volume, and updating a physical volume's status information based upon a combination of that volume's prior status and that volume's current status; specifying a time duration that expired logical unit data may remain on a physical volume; based upon the volume's updated status information, determining whether a selected volume has expired data longer than the specified time duration; and, overwriting contents of the volume having data that has expired longer than the specified time duration. (See, e.g., Abstract and paragraphs 15-20.) However, unlike the present invention, Gibble et al. does not disclose, at a minimum, sequential procedures of: holding a first table that, if files as backup objects are classified into plural groups, stores each group name file names, and each policy in a manner of being made to correspond to each other, said file names being included in said each group, said each policy being a method for managing replication data. Furthermore, Gibble et al. does not disclose when transmitting pair information to a storage apparatus in response to a receipt of a specified second group name, a second group including a first file, a pair information including a first original-volume name and a second duplicate-volume name, a first original-volume storing therein the data of said first file, extracting a first duplicate-volume name out of a second table, said first duplicate-volume storing therein replication data of said first file. More particularly, Gibble et al. does not disclose or suggest the above described first feature of the present invention as recited in independent

claim 1, and the above described second feature of the present invention as recited in independent claim 2, in combination with the other limitations recited in each of the independent claims.

U.S. Patent Publication No. 2004/0059882 (Kedem et al.) discloses a method and system for copying data from one storage device to another storage device. In response to a copy command from a requesting host application identifying the storage locations in source and destination disk storage devices, an environment is established. Thereafter a calling system receives an immediate response that the copy operation is complete even though no data has been copied. Application programs may access storage locations in either the source or the destination disk storage device during the copy phase. A copy program transfers the date on a track-by-track basis to the storage locations in the destination disk storage device. Procedures assure that any data access to a particular track in either the source or destination disk storage devices by any application prior to the transfer of that track are accommodated to maintain data integrity. The system utilizes a flag that can be set to either allow or disallow the overwriting of a copied file track. (See, e.g., Abstract and paragraphs 13-15, and 69.) However, unlike the present invention, Kedem et al. does not disclose, at a minimum, sequential procedures of: holding a first table that, if files as backup objects are classified into plural groups, stores each group name file names, and each policy in a manner of being made to correspond to each other, said file names being included in said each group, said each policy being a method for managing replication data. Furthermore, Kedem et al. does not disclose when


transmitting pair information to a storage apparatus in response to a receipt of a specified second group name, a second group including a first file, a pair information including a first original-volume name and a second duplicate-volume name, a first original-volume storing therein the data of said first file, extracting a first duplicate-volume name out of a second table, said first duplicate-volume storing therein replication data of said first file. More particularly, Kedem et al. does not disclose or suggest the above described first feature of the present invention as recited in independent claim 1, and the above described second feature of the present invention as recited in independent claim 2, in combination with the other limitations recited in each of the independent claims.

Therefore, since the references fail to disclose or suggest the above described first feature of the present invention as recited in independent claim 1, and the above described second feature of the present invention as recited in independent claim 2, in combination with the other limitations recited in each of the independent claims, it is submitted that all of the claims are patentable over the cited references.

In view of the foregoing, Applicant requests that this Petition to Make Special be granted and that the application undergo the accelerated examination procedure set forth in MPEP 708.02 VIII.

Respectfully submitted,

MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.



Frederick D. Bailey
Registration No. 42,282

FDB/sdb
(703) 684-1120